Journal of Digital Imaging

Automated Recognition of Lateral from PA Chest Radiographs: Saving Seconds in a PACS Environment

John M. Boone, Ph.D., Greg S. Hurlock, M.D., J. Anthony Seibert, Ph.D., and Richard L. Kennedy, M.S.

Images acquired in a two-view digital chest examination are frequently not electronically distinguishable. As a result the lateral and posterioanterio (PA) images are often improperly positioned on a PACS work station. A series of 1998 chest radiographs (999 lateral, 999 PA or AP) were used to develop a neural network classifier. The images were down-sampled to 16 × 16 matrices, and a feed-forward neural network was trained and tested using the "leave-one-out" method. Using five nodes in the hidden layer, the neural network correctly identified 987 of the 999 test cases (98.8%) (average of six runs). The simple architecture and speed of this technique suggests that it would be a useful addition to PACS work station software. The accumulated time saved by correctly positioning the lateral and PA chest images on the work station monitors in accordance with each radiologist's hanging protocols was estimated to be about 1 week of radiologist time per year.

KEY WORDS: Chest radiography, picture archiving and communication system (PACS), pattern recognition, neural networks

THE USE OF A PICTURE archiving and communication system (PACS) is becoming necessary in the modern radiology department.[1-3] The fledgling PACS systems of the 1980s have evolved, in general, into high-throughput tools that allow both primary diagnosis and secondary viewing. Digital chest radiographs make up a substantial fraction (~40%) of the radiographic images read out on a PACS system. Computed radiography[4-6] (CR) is now the most widely used digital image detector for digital radiography, and is likely to remain so for some time, especially in the portable setting. The chest radiographic examination typically is comprised of both the lateral and posteroanterior (PA) views for upright examinations and the lateral and antereoposterio (AP) views in the portable environment. Subsequent reference will be to the PA view; however, for the purposes of this study, this should be considered interchangeable with the AP view. The two cassettes used per patient are typically not tracked by the technologist in terms of which one is the lateral and which one is the PA image. Consequently, when the images are read out in a CR reader, they are named with the same generic label (e.g., "chest PA/Lat"). Neither the PACS system nor its user can uniquely identify which image file is the PA and which is the lateral.

Because each radiologist has preferences in terms of hanging protocol, it would be useful for the PACS system to know unambiguously which image is the PA and which is the lateral, for each patient's pair of chest radiographs. The purpose of this investigation was to develop pattern-recognition techniques that can distinguish between the PA and the lateral views in a pair of chest radiographs for a given patient. This identification step will lead to higher throughput by enabling the PACS system to correctly "hang" the PA and the lateral images in accordance with the preferences of the radiologist.

Although switching a pair of images on a high-resolution work station only requires a few moments, given the large number of chest radiographs reviewed, the accumulated time spent

From the Department of Radiology, University of California, Davis Medical Center, Sacramento, CA, 95817, USA.

Correspondence to: John M. Boone, Ph.D., tel: 916-734-3158; fax: 916-734-0316; e-mail: jmboone@ucdavis.edu

Online publication 30 January 2004

doi: 10.1007/s10278-003-1662-y

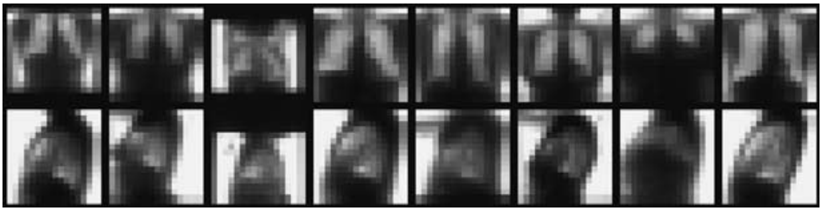

Fig 1. A collection of eight pairs of small (16 × 16 pixels) images are illustrated. These individual images represent the inputs to the neural network described in the text.

can be substantial. Moreover, there is an annoyance factor associated with the need to swap images that can be largely eliminated with the proposed technique.

MATERIALS AND METHODS

A neural network approach was used to determine which image is the PA and which is the lateral of a pair of chest radiographs. Neural network techniques[7-9] are relatively simple, have been used extensively in medical image pattern-recognition with success, and are known to a large number of investigators familiar with neural network implementation.

A series of 1,000 consecutively selected chest radiograph pairs were downloaded from the clinical PACS system at our institution under an approved IRB protocol. Custom software was written on an imaging workstation (Microsoft Windows 2000 and Microsoft Visual C/C++ 5.0, Redmond, WA) to manipulate and reformat the image data. Initial software was written which allowed a trained observer (fourth-year medical student at the time) to indicate which of the two images (per patient) was the PA and which was the lateral. One image pair was corrupted (both images were PA), and thus this pair was eliminated from the data set, leaving 999 image pairs for a total of 1,998 individual images. Each original image was acquired with a pixel matrix of 1760 × 2140, corresponding to ~200 μm pixels. The rectangular images were reduced in resolution to a 16 × 16 square image, where each pixel value in this much smaller image represented the mean gray scale value of a 110 × 134 pixel rectangular region in the original high-resolution clinical image. The 256 (16 × 16) pixels representing this small image were rescaled to the range {0.0, 1.0}. These 256 resulting values were used as the input data to a feed-forward artificial neural network.[7] Examples of the image data set used are shown in Figure 1.

Neural network software developed by one of the authors and reported previously[7] was used in this study. The "jackknife" or "leave-one-out" method was used for training and testing. For N images (N = 1,998 here), one image was removed from the data set, the remaining N −1 images were used to train the neural network, and that trained network was used to evaluate the removed image. This process was repeated until each image had a turn being left out from the training procedure and being used for performance assessment. Note that the left-out image does not participate in the training, and therefore each of the individual jackknife evaluations was performed on an independent image. Thus, the evaluations performed in this study included the well-known effects of "shrinkage."[10] Shrinkage refers to the reduction in performance when an independent testing data set is used, compared to the performance on the data set used for training. The jackknife approach allowed the evaluation of a test set, independent of each training set, and thus the influence of shrinkage is included in our performance analysis.

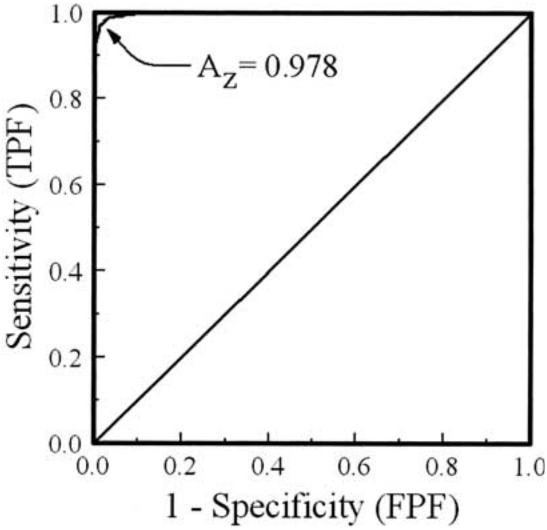

Fig 2. A receiver operating characteristic (ROC) curve is illustrated. The classification performance of the five-node hidden layer network demonstrated excellent performance, with an area under the ROC curve (A_z) of 0.978.

A two-layer neural network (one hidden layer) was used. It was comprised of 256 nodes in the input layer, one to nine nodes in the hidden layer, and one node in the output layer. A sigmoid activation function was used at each node,[7] and each node was individually biased. All weights of the neural network were initialized to a range of {−0.30, +0.30} using a random number generator. A momentum term of 0.700 and a learning rate coefficient of 0.100 were used.[7] The 1,998 training and testing sessions were performed on a 2.5 GHz Pentium-based (Intel Corporation, Sunnyvale, CA) work station.

The number of hidden nodes necessary to perform accurate classification was studied by evaluating the performance over a range of hidden node values from 1 to 9. Each jackknife run of 999 cases was run six times for networks with from 1 to 6 nodes. The 7-node architecture was evaluated five times, and 8- and 9-node architectures were run four times each.

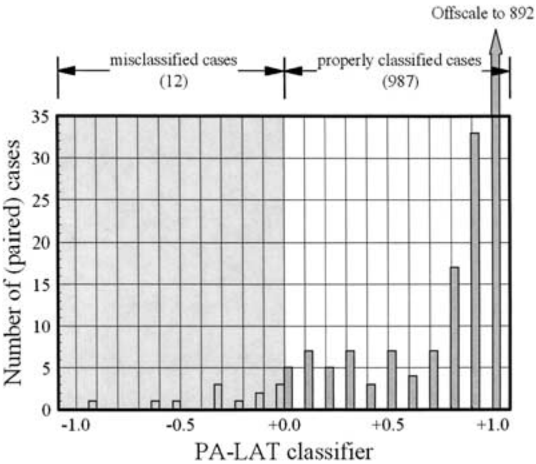

Fig 3. The number of cases as a function of the PA-lat classifier is illustrated. Negative values of the classifier correspond to misclassified cases, and positive values correspond to properly classified cases. Notice that the vast majority of cases achieved near-perfect classification at +1.0, where the height of that histogram bar is off-scale.

The pattern of weights employed by the neural network for identifying the PA from the lateral chest images was studied with a neural network with one hidden node. Although the overall performance of a one-node hidden layer network was lower than that for larger hidden layers, the analysis of weights is more straightforward in the situation where each of the input nodes (from the image) has only one weight affiliated with it.

For each image, the neural network outputs a raw score (call this α) that ranges from 0 to 1 (inclusive). Perfect identification corresponds to assigning α to 1 for PA images, and α to 0 for lateral images. Let α_{PA} and α_{LAT} correspond to the neural network's assignment for a given patient's pair of chest images. A classification parameter β can be defined such that:

$$\beta = \alpha_{\mathrm{PA}} - \alpha_{\mathrm{LAT}}$$

For proper classification we require that $\alpha_{PA} > \alpha_{LAT}$, and this condition is met when $\beta > 0.0$. Negative values of β therefore correspond to improperly classified lateral and PA images.

RESULTS

The neural network that used five hidden nodes demonstrated 98.8% accuracy, misclassifying 12 cases (the average of six runs was 12.2 misclassified cases), but properly classifying 987 cases. Figure 2 shows a receiver operating characteristic (ROC) curve that demonstrates the excellent performance of the neural network-based PA/LAT classifier.

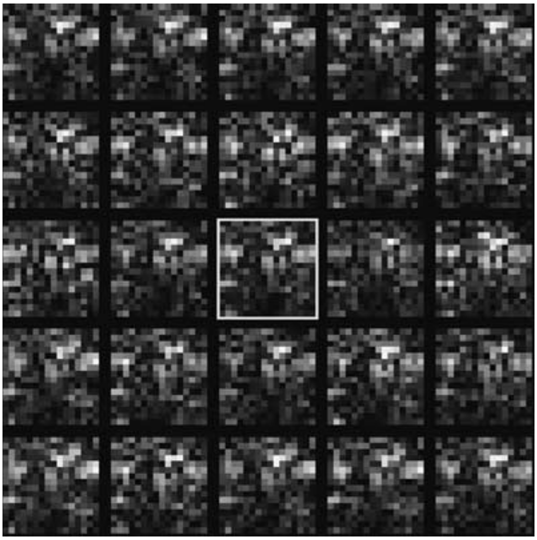

Fig 4. A patchwork of 24 arbitrarily selected neural network weight patterns is illustrated, and the centermost pattern (outlined in white) corresponds to the average weight pattern for all 1998 training sets. The weight patterns here correspond to a 1 node hidden layer. Lighter values correspond to positive and darker values correspond to negative weight values. These weight patterns illustrate the location of features that are used by the neural network for distinguishing between the PA and lateral chest radiographic images.

The distribution of cases as a function of the classification parameter β is shown in Figure 3, for the five-node hidden layer neural network architecture. The value of β runs from -1 to $+1$, and positive values correspond to correct identification of the PA and lateral chest images. Nearly 90% of the β values are near perfect (0.9-1.0), and 96% of these values are greater than 0.5. The 12 pairs of images that were misclassified correspond to cases where β was negative. The images of these 12 cases were visually evaluated; however, they did not differ in any obvious manner from the remainder of the images in the data set.

Figure 4 demonstrates the distribution of weights for the one hidden layer neural network. This network achieved 97.6% (974/999) correct classification performance (averaged over six runs), slightly lower than the five hidden node network performance. The one hidden node network is more useful in understanding the spatial distribution of weights used by the neural network classifier, because with this network architecture only one weight

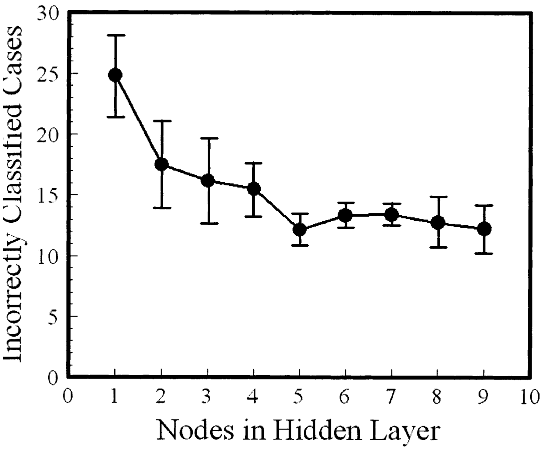

Fig 5. The performance of the neural network classification as a function of the number of nodes in the hidden layer is demonstrated. Increasing the number of hidden nodes improved testing performance, although diminishing returns are seen beyond five nodes.

value is assigned per pixel to the small 16 × 16 images that served as inputs to the neural network. A collection of 25 weight distributions (selected randomly from the 1,998 weight arrays) is shown in Figure 4. The center weight pattern (outlined in white in Figure 4) demonstrates the ensemble average of all of the 1,998 weight patterns of one jackknife run. Figure 4 illustrates that much of the positive weighting (brighter regions in Figure 4) is located approximately at the point where the arms are situated on the PA view, and where the body contour transitions to air on the lateral image (see Fig. 1). There is also a prominent region of negative weighting (shown as dark values in Figure 4) near the bottom center of the images, corresponding to the position of the mediastinum on the PA images.

Figure 5 illustrates the number of misclassified cases as a function of the number of nodes in the hidden layer. The mean is illustrated with error bars indicating $\pm 1\sigma$, and these values were assessed from multiple jackknife runs for each node, as mentioned in Materials and Methods. The computational dexterity of a neural network increases as the number of hidden nodes increases, and this flexibility is shown to be advantageous for five or more hidden nodes.

DISCUSSION

The goal of this project was to develop an automatic technique that could be used by a PACS work station to correctly "hang" the PA and lateral chest images for radiologist viewing and interpretation. Unlike computer-aided diagnostic (CAD) tools,[11,12] the penalty for misclassification in this application is simply a couple of extra seconds spent swapping the PA and the lateral images on a PACS work station, not an incorrect diagnosis. Given this, it is not necessary to achieve 100% performance, and the approximately 99% performance demonstrated by the neural network system developed here is considered adequate.

How would this simple classifier system affect PACS efficiency? A large institution may produce on the order of 200 two-view chest radiographs in a 24-hour period, and this corresponds to about 73,000 two-view chest examinations per year. Assume that swapping incorrectly hung chest images requires 4 seconds of radiologist time. A PACS system will hang approximately 50% of cases correctly in the absence of guidance (by simply guessing), and the neural network discussed here can achieve about 99% correct performance. Using these numbers, 40.5 hours per year would be spent swapping chest images in the absence of guidance, and 48.6 minutes per year would be spent with the neural network guidance. This leads to a net savings of about one week (39.7 hours) of radiologist time per year. At a rate of $150 per hour ($312,000 per year), the neural network classifier described in this study would save about $6,000 per year. In addition to monetary savings, a reduction in radiologist frustration may result as well.

The implementation of the proposed technique requires only modest effort, but it lies in the hands of those who develop PACS systems. To facilitate adoption of the technique, the weight values and mathematical necessities for their use will be made available to any interested party. The development of a PACS system requires that attention be paid to a large number of issues. Although the classification technique described here is a relatively minor aspect of a PACS, it is the implementation of a large number of small features such as the one reported here that adds value to state-of-the-art PACS systems.

REFERENCES

1. Siegel, E, Reiner, B: Work flow redesign: the key to success when using PACS. AJR Am J Roentgenol 178:563-566, 2002

2. Passadore, DJ, Isoardi, RA, Ariza, PP, et al: Use of a low-cost, PC-based image review workstation at a radiology department. J Digit Imaging 14:222-223, 2001

3. Twair, AA, Torreggiani, WC, Mahmud, SM, et al: Significant savings in radiologic report turnaround time after implementation of a complete picture archiving and communication system (PACS). J Digit Imaging 13:175-177, 2000

4. Nakano, Y, Gido, T, Honda, S, et al: Improved computed radiography image quality from a BaFl:Eu photostimulable phosphor plate. Med Phys 29:592-597, 2002

5. Aoyama, M, Li, Q, Katsuragawa, S, et al: Automated computerized scheme for distinction between benign and malignant solitary pulmonary nodules on chest images. Med Phys 29:701-708, 2002

6. Freedman, M: Chest image quality: film-screen vs. computed radiography. Adm Radiol 10:38-40, 42, 1991

7. Boone, JM, Sigillito, VG, Shaber, GS: Neural networks in radiology: an introduction and evaluation in a signal detection task. Med Phys 17:234-241, 1990

8. Boone, JM, Gross, GW, Greco-Hunt, V: Neural networks in radiologic diagnosis. I. Introduction and illustration. Invest Radiol 25:1012-1016, 1990

9. Gross, GW, Boone, JM, Greco-Hunt, V, et al: Neural networks in radiologic diagnosis. II. Interpretation of neonatal chest radiographs. Invest Radiol 25:1017-1023, 1990

10. Boone, JM: Sidetracked at the crossroads. Radiology 193:28-30, 1994

11. Li, Q, Katsuragawa, S, Doi, K: Computer-aided diagnostic scheme for lung nodule detection in digital chest radiographs by use of a multiple-template matching technique. Med Phys 28:2070-2076, 2001

12. Doi, K, Giger, ML, Nishikawa, RM, et al: Computer-aided diagnosis of breast cancer on mammograms. Breast Cancer 4:228-233, 1997